Journal of Engineering and Multidisciplinary Applied Sciences
Vol. 01, Issue 01, July–December 2026, pp. 45-54

On-Device Detection Is Not an Optimisation but a Precondition: Energy Budget Analysis of a Cellular-Connected Ambulatory Cardiac Monitor

Mr. Pranesh Balaso Bamankar

Assistant Professor, Mechanical Engineering, Sharadchandra Pawar College of Engineering, Someshwarnagar, Baramati, Maharashtra, India.

Email: praneshmtech@gmail.com

ABSTRACT

Long-term ambulatory cardiac monitors are increasingly built as standalone connected devices that upload over a cellular network rather than through a paired smartphone. That architectural choice changes the energy economics completely, because a cellular session costs orders of magnitude more energy than a short-range link, and it determines whether raw signal streaming is available to the designer at all. This study quantifies the average current and achievable recording duration of four upload strategies for a cellular-connected single-lead monitor, identifies which subsystem limits each, and establishes how the design is affected by variables the engineer does not control — the patient's event rate and the detector's specificity. An energy budget model was constructed for a 500 mAh device sampling one electrocardiographic lead at 256 Hz and 12 bits. Subsystem currents were 0.35 mA for the analogue front end, 0.25 mA for the buffering microcontroller, 4.0 mA while the detection routine runs, 180 mA during cellular transmission and 45 mA during reception, with a 2.2 s attach overhead and 0.55 s per kilobyte transmitted. Four strategies were compared: continuous raw upload, hourly summary without detection, on-device detection with 30 s event strips, and the same with fourfold strip compression. Event rate was swept from 2 to 200 per day and detector specificity from 99.99 % to 99.5 %. Continuous raw upload required 368.5 mA average current and gave 0.051 days of recording — 73 minutes — with the radio accounting for 99.8 % of the budget. On-device detection with 12 event strips per day required 1.086 mA and gave 17.3 days; adding fourfold compression gave 18.7 days. Once detection was moved on-device the energy budget became balanced rather than radio-limited: the analogue front end contributed 32.2 %, the idle microcontroller 23.0 %, the detection routine 22.1 %, the radio 21.5 % and sleep 1.1 %. Recording duration fell from 20.7 days at 2 events per day to 4.2 days at 200. Degrading detector specificity from 99.99 % to 99.5 % generated five false strips per day and cost 1.3 days of recording. In a cellular-connected monitor, on-device detection is not an efficiency measure but the condition under which the device functions at all, since streaming is short by a factor of about 340. Once it is adopted no single subsystem dominates, so further gains require simultaneous work on four fronts, and the largest remaining determinant of recording duration is the patient's event rate.

Keywords: wearable biomedical device; edge computing; energy budget; cellular IoT; ambulatory monitoring; embedded design.

1. INTRODUCTION

Ambulatory cardiac monitoring has moved from a recorder that is worn, returned and analysed offline toward a connected device that reports while it is still on the patient [1-3]. The

clinical motivation is straightforward: an arrhythmia detected on day two is more useful than the same arrhythmia found on day fifteen when the device is posted back.

Connectivity can be provided in two ways, and the choice determines the engineering problem. A short-range link to a paired smartphone is cheap in energy terms and expensive in every other respect [14]: it requires the patient to own, carry and charge a compatible phone, to keep an application running, and to remain within range, and adherence to all three is imperfect in exactly the elderly population most often monitored. A direct cellular link removes those dependencies and replaces them with an energy problem, because a cellular transmission draws roughly two orders of magnitude more current than a short-range one and carries a fixed attach overhead that is paid whether the payload is one byte or one kilobyte [9-11,15].

The question this creates for the designer is what the device should send. Raw signal offers the greatest clinical flexibility, since analysis can be revised later and the full record is available for adjudication. Extracted features are compact but discard the evidence a clinician needs to confirm a diagnosis. Detected event strips are a compromise that requires the device itself to decide what is worth sending, which places a diagnostic algorithm inside a battery-powered wearable [6-8,13].

This paper quantifies the trade. Three questions are addressed: what recording duration each strategy actually permits; which subsystem limits the design in each case; and how sensitive the result is to two variables the engineer does not control, namely how often the patient has events and how often the detector is wrong.

2. SYSTEM MODEL

Figure 1 shows the signal chain of the modelled device and the current drawn by each stage.

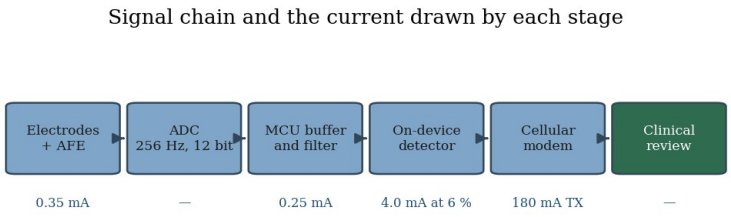


Fig.1: Signal chain of the modelled device, with the current drawn by each stage.

2.1 Device and Subsystem Currents

The modelled device is a single-lead adhesive patch monitor with a 500 mAh lithium cell, of which 90 % is treated as usable before the cut-off voltage [2,3]. One electrocardiographic channel is sampled continuously at 256 Hz with 12-bit resolution, giving a raw data rate of 384 bytes per second. The subsystem parameters are listed in Table 1.

2.2 Energy Accounting

The charge consumed by one cellular session is the sum of the attach overhead, the payload transmission and the acknowledgement [10,12]:

Table 1: Subsystem parameters.

Subsystem	Current	Duty	Note
Analogue front end and bias	0.35 mA	Continuous	Cannot be duty-cycled without losing signal
Microcontroller, buffering and filtering	0.25 mA	Continuous	Low-power mode between samples
Microcontroller, detection routine	4.00 mA	6 % in strategies C and D	Filter bank, QRS detection, rhythm classification
Cellular modem, transmit	180 mA	Per session	Dominates any session
Cellular modem, receive	45 mA	1.0 s per session	Acknowledgement and control
Attach and connection overhead	180 mA	2.2 s per session	Paid once per session regardless of payload
Transmission time	—	0.55 s per kilobyte	Marginal cost of payload
Sleep, power-saving mode	0.012 mA	Remainder	Between sessions
Battery	500 mAh, 90 % usable	—	Adhesive patch form factor

Note: Values are representative of contemporary low-power analogue front ends, Cortex-class microcontrollers and NB-IoT or LTE-M modems [4,5,10], and are stated assumptions rather than measurements of any product. Absolute durations are specific to them; the comparisons between strategies follow from the structure of the budget.

$$Q_{session} = I_{TX}\left(t_{attach} + t_{kB} \times P\right) + I_{RX}\, t_{RX} \tag{1}$$

with P the payload in kilobytes. Average device current is then

$$I_{avg} = I_{AFE} + I_{MCU} + I_{DSP} \times d + \frac{N\, Q_{session}}{86400} + I_{sleep} \tag{2}$$

with N sessions per day and d the detection duty cycle. The structure of Equation (1) is the reason session count matters more than payload size: the 2.2 s attach costs 396 mA·s whether the device sends 100 bytes or 1 kB, so a strategy that sends little but often is penalised heavily.

2.3 Strategies Compared

Table 2 lists the four upload strategies compared.

Table 2: Upload strategies.

Strategy	What is uploaded	Sessions per day	Detection on device?
A: continuous raw upload	Every sample, in 512-byte payloads	64,800	No
B: hourly summary	Aggregated rate and interval statistics, 1 kB	24	No
C: on-device detection	30 s raw strip per detected event, plus a daily summary	13	Yes
D: detection with compression	The same strips at fourfold compression	13	Yes

3. RESULTS

3.1 Recording Duration

Table 3 and Figure 2 give the average current and recording duration under each strategy.

Table 3: Energy and duration by strategy.

Strategy	Sessions/ day	Uploaded per day	Average current	Radio share	Recording duration
A: continuous raw upload	64,800	32.4 MB	368.5 mA	99.8 %	0.051 days (73 min)
B: hourly summary	24	24 kB	0.762 mA	19.7 %	24.6 days
C: on-device detection	13	136 kB	1.086 mA	21.5 %	17.3 days
D: detection + compression	13	34.8 kB	1.000 mA	10.8 %	18.7 days

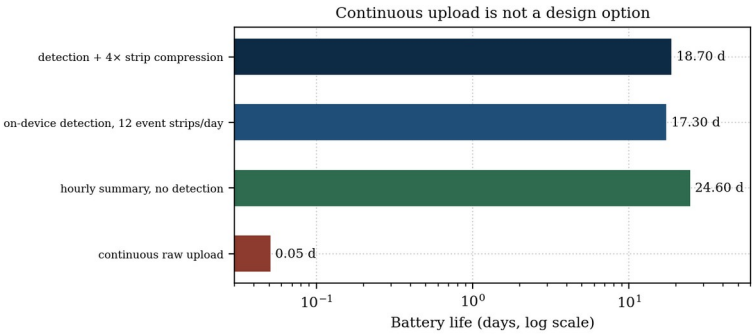


Fig.2: Recording duration by strategy, on a logarithmic scale.

Continuous raw upload gives 73 minutes of recording. That is not a poor design to be improved but an unavailable one: the device would need a battery 340 times larger to reach a fourteen-day wear period, which for an adhesive patch is not a packaging problem but an impossibility.

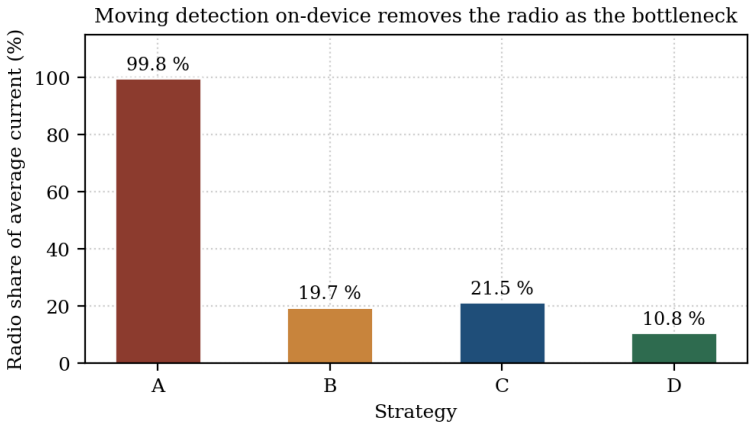


Fig.3: Share of average current drawn by the cellular radio under each strategy.

This is the finding that motivates the paper's title. In a smartphone-tethered architecture, streaming raw signal is entirely practical and on-device analysis is an optimisation that trades microcontroller cycles for radio time. In a cellular architecture the same choice is not an

optimisation, because one of the options does not exist. Any cellular-connected long-wear monitor has a detection algorithm inside it, whether or not its designers describe the device that way.

3.2 What Limits Each Design

As Figure 3 shows, the radio accounts for 99.8 % of the budget under continuous upload and between 10.8 % and 21.5 % once the device decides for itself what to send. The transition is abrupt rather than gradual, because the number of sessions falls from 64,800 to 13 and each session carries a fixed attach cost.

Table 4: Energy budget of strategy C.

Subsystem	Share of average current
Analogue front end	32.2 %
Microcontroller, idle and buffering	23.0 %
Detection routine	22.1 %
Cellular radio	21.5 %
Sleep	1.1 %

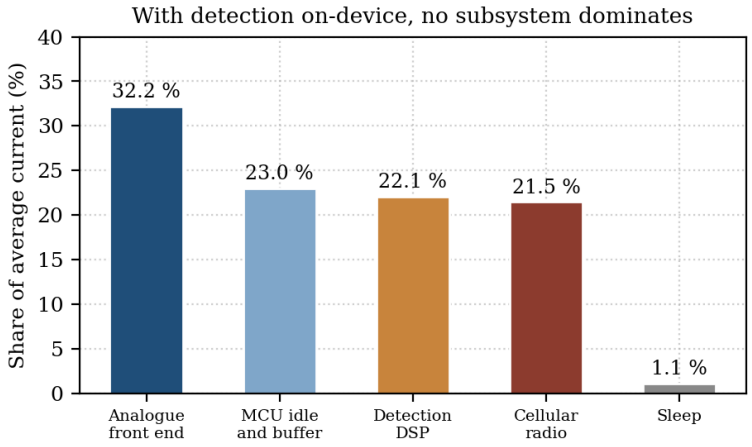


Fig.4: Energy budget of the on-device detection design.

Table 4 and Figure 4 give the more useful engineering result, and it is easy to miss. Once detection is on the device, the budget is almost flat: four subsystems contribute between 21.5 % and 32.2 % each. There is no bottleneck, and therefore no single change that materially extends the recording.

The practical consequence is a change in what optimisation is worth attempting. Halving the radio energy — through compression, better scheduling or a more efficient modem [16] — extends the recording by about 12 %, which is what strategy D achieves at 18.7 days against 17.3. Halving the detection routine's cost gains a similar amount [17]. Neither is transformative, and a design effort aimed at any one of them will produce a disappointing result that is easy to misattribute to poor execution.

The analogue front end is the largest single item at 32.2 %, and it is the one that cannot be duty-cycled at all, because an electrocardiographic channel that is switched off is not recording [4,5]. It is also the item least often targeted in power optimisation work, which tends to concentrate on the radio for the good reason that the radio is dominant in most other IoT applications.

3.3 Event Rate: The Variable the Designer Does Not Control

Table 5 and Figure 5 give the recording duration as the number of detected events per day increases.

Table 5: Recording duration against detected event rate.

Events per day	2	6	12	24	48	96	200
Recording duration (days)	20.7	19.2	17.3	14.4	10.8	7.2	4.2

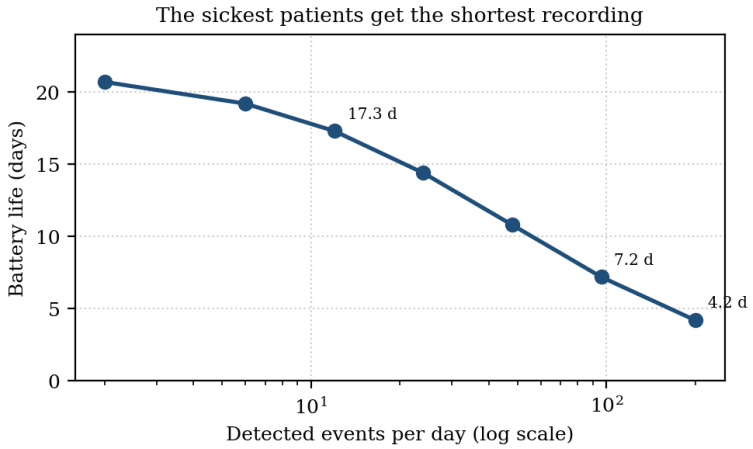


Fig.5: Recording duration against the number of events detected per day.

Recording duration falls from 20.7 days at two events per day to 4.2 days at two hundred. The relationship is a direct consequence of Equation (1): each event triggers a session and each session costs the attach overhead.

The clinical implication is uncomfortable and deserves to be stated plainly. A patient with rare, isolated ectopy records for three weeks; a patient in frequent paroxysmal atrial fibrillation with many detected episodes records for less than a week. The device gives the shortest recording to the patients who have the most to detect, and the prescribing clinician has no way of knowing in advance which patient they are dealing with — that is, after all, why the monitor was fitted.

Several mitigations follow from the same arithmetic. Batching events and uploading them on a schedule rather than immediately spreads the attach overhead across multiple strips and removes most of the dependence, at the cost of detection latency. Capping strips per day preserves battery and loses evidence. Adapting the strip length or the detection threshold once a high event rate has been observed is a form of closed-loop behaviour that raises its own regulatory questions. None is free, and the choice between them is a clinical decision presented as a firmware one.

3.4 Detector Specificity

Table 6 and Figure 6 give the effect of detector specificity on recording duration.

Degrading specificity from 99.99 % to 99.5 % produces five false strips per day and costs 1.3 days of recording, a 7.6 % reduction. The effect is real but secondary to the event rate, which is reassuring for the algorithm designer [6,7].

It is worth noting because the usual argument for detector specificity concerns the workload it creates for whoever reviews the strips, and that argument is often met with the

response that reviewing a few extra traces is cheap. On this model the false positives also consume battery, and battery is what determines whether the monitor is still recording on the day the diagnostic event occurs. Specificity is therefore a device performance parameter as well as a workflow one, and the two arguments should be made together.

Table 6: Effect of detector specificity.

Specificity	False strips per day	Recording duration	Loss against best case
99.99 %	0.1	17.2 days	—
99.95 %	0.5	17.1 days	0.1 days
99.90 %	1.0	17.0 days	0.2 days
99.50 %	5.0	15.9 days	1.3 days

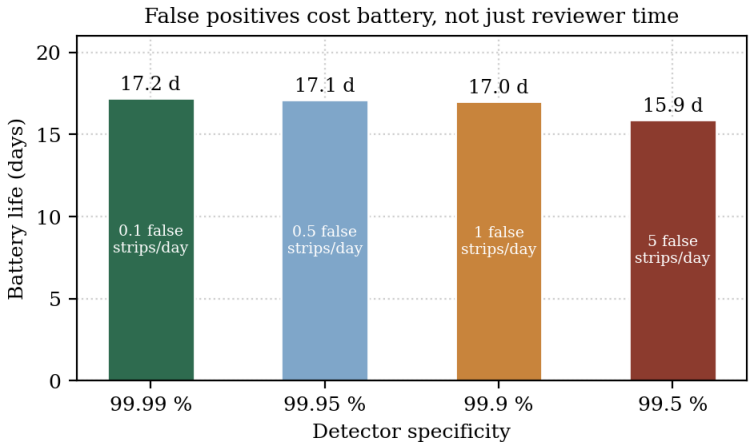


Fig.6: Recording duration against detector specificity.

4. DISCUSSION

4.1 Principal Findings

Four results follow. Continuous raw upload over a cellular link gives 73 minutes of recording against a fourteen-day requirement, so it is unavailable rather than inefficient. On-device detection gives 17.3 days and compression 18.7. Once detection is on the device the budget is flat, with four subsystems each contributing between 21.5 % and 32.2 %, so no single optimisation is decisive. And the strongest remaining determinant of recording duration is the patient's event rate, which falls from 20.7 to 4.2 days across the range examined.

4.2 Design Implications

Table 7 summarises each finding, its implication and the corresponding design rule.

Table 7: Implications.

Finding	Implication	Design rule
Streaming is short by a factor of 340	Architecture choice forecloses the option	Treat on-device detection as a requirement of cellular connectivity
Attach overhead is paid per session	Session count matters more than payload	Batch uploads; do not send each event as it occurs unless latency demands it

Finding	Implication	Design rule
Budget is flat once detection is local	No single bottleneck remains	Plan simultaneous gains, or accept the duration the architecture gives
Front end is the largest item and cannot be duty-cycled	The usual radio-first optimisation is misdirected	Specify the analogue front end as carefully as the modem
Compression bought 8 %	Payload reduction has limited headroom	Do not expect codec work to change the product category
Duration falls 5-fold across event rate	The sickest patients record shortest	Design the upload schedule to be event-rate insensitive
False positives cost 1.3 days at 99.5 % specificity	Specificity is a hardware parameter too	Include battery cost in the detector operating-point decision

The sixth row is the one with the clearest product consequence. A monitor whose advertised wear time is fourteen days will meet that claim for a patient with few events and fall well short for a patient with many, and the discrepancy will appear in the field rather than in bench testing, because bench testing is usually done at a nominal event rate. Specifying wear time at a stated event rate, or designing a schedule that decouples the two, would prevent a class of field failure that looks like a manufacturing defect and is a specification omission.

The fourth row is a corrective to a habit. In most connected sensing applications the radio dominates the budget so heavily that power optimisation is effectively radio optimisation, and that intuition transfers poorly here. In a device that must sense continuously to be useful, the always-on analogue path becomes the largest consumer once the radio duty cycle collapses, and it is the one component whose consumption cannot be reduced by sending less.

4.3 Limitations

The model is an energy budget, not a measurement. Subsystem currents are representative values for the component classes named and would differ between specific parts by a factor that matters; modem current in particular varies substantially with signal strength, and a device at the edge of coverage pays far more per session than the model assumes, which would worsen every strategy but worsen the session-heavy ones most [10,11]. Battery capacity is treated as constant, whereas usable capacity falls with temperature, age and discharge rate, so the durations reported are optimistic.

The detection routine is characterised by a single duty cycle and current rather than by an implementation. A different algorithm, a hardware accelerator or a lower clock would change the 22.1 % contribution substantially [8,17], and the finding that the budget is flat is specific to the duty cycle assumed; with a much cheaper detector the front end would dominate and with a much more expensive one the design would be compute-limited. That sensitivity is the model's main structural weakness and could be resolved only by profiling a real implementation.

Several real behaviours are omitted. Signal quality is assumed constant, whereas electrode contact degrades over a wear period and generates both artefact and false detections; the model has no mechanism for retransmission after a failed session, which is common on cellular links and adds energy unpredictably; and no allowance is made for firmware update, time synchronisation or periodic keep-alive traffic, all of which consume sessions. Each of these makes the modelled durations optimistic rather than conservative.

Finally, the analysis addresses energy and says nothing about diagnostic performance. A device that records for eighteen days while missing the arrhythmia is worse than one that records for seven and finds it, and the relationship between on-device detector sensitivity and diagnostic yield is a separate question that this model does not address and that would determine whether the strategies compared here are clinically equivalent.

5. CONCLUSION

For a cellular-connected single-lead cardiac monitor, continuous raw upload draws 368.5 mA and gives 73 minutes of recording, against 1.086 mA and 17.3 days for a device that detects events itself and uploads only strips. The ratio is about 340, which places streaming outside the design space entirely rather than at its inefficient end. On-device detection is therefore a precondition of the architecture and not a refinement of it.

Having adopted it, the designer inherits a flat energy budget in which the analogue front end, the idle microcontroller, the detection routine and the radio each account for roughly a fifth to a third of consumption. There is no bottleneck left to remove, so incremental effort returns incremental gains: fourfold payload compression bought 8 % of recording time. The habit of treating power optimisation as radio optimisation, which serves most connected sensing well, is actively misleading here, because the always-on sensing path is the single largest consumer and is the one that cannot be made cheaper by sending less.

The largest remaining lever is not a component at all. Recording duration fell from 20.7 days to 4.2 across the range of event rates examined, so the patient determines the wear time more than the hardware does, and a device specified at a nominal event rate will underperform in exactly the patients it was fitted to investigate. Designing the upload schedule to be insensitive to event rate is worth more than any component substitution considered here.

DECLARATIONS

Ethical approval. Not applicable. This is a computational design study; no human participants, patient data, clinical recordings or physical devices were involved.

Funding. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Conflicts of interest. The author declares no conflicts of interest. The author has no financial or advisory relationship with any device manufacturer, semiconductor supplier or network operator.

Product and regulatory disclaimer. No commercial device, component, modem or network is evaluated, and no result describes any marketed product. Nothing here constitutes a claim of clinical performance or fitness for purpose; an ambulatory monitor is a regulated medical device and its wear time, detection performance and safety are established by validation, not by an energy model.

Parameter disclaimer. All subsystem currents, timings and battery values are representative assumptions for the component classes named. Absolute durations apply only to the case modelled; the comparisons follow from the structure of the budget, as set out in Section 4.3.

Data and code availability. This is a computational design study; no measurements, patient recordings or external datasets were used. All results are derived from Equations (1) and (2) using the parameters listed in Table 1, and every reported value can be reproduced from them. The analysis code is available from the corresponding author upon reasonable request.

Author's contribution. Pranesh Balaso Bamankar is the sole author of this work and was responsible for the conception and design of the study, development of the energy budget model, analysis and interpretation of the results, preparation of the figures and tables, and drafting, revision and final approval of the manuscript. The author is accountable for all aspects of the work.

REFERENCES

1. Steinhubl SR, Waalen J, Edwards AM, et al. Effect of a home-based wearable continuous ECG monitoring patch on detection of undiagnosed atrial fibrillation: the mSToPS randomized clinical trial. JAMA. 2018;320(2):146–155.
2. Barrett PM, Komatireddy R, Haaser S, et al. Comparison of 24-hour Holter monitoring with 14-day novel adhesive patch electrocardiographic monitoring. Am J Med. 2014;127(1):95.e11–95.e17.
3. Turakhia MP, Hoang DD, Zimetbaum P, et al. Diagnostic utility of a novel leadless arrhythmia monitoring device. Am J Cardiol. 2013;112(4):520–524.
4. Yazicioglu RF, Kim S, Torfs T, Kim H, Van Hoof C. A 30 µW analog signal processor ASIC for portable biopotential signal monitoring. IEEE J Solid-State Circuits. 2011;46(1):209–223.
5. Van Helleputte N, Konijnenburg M, Pettine J, et al. A 345 µW multi-sensor biomedical SoC with bio-impedance, 3-channel ECG, motion artifact reduction, and integrated DSP. IEEE J Solid-State Circuits. 2015;50(1):230–244.
6. Pan J, Tompkins WJ. A real-time QRS detection algorithm. IEEE Trans Biomed Eng. 1985;BME-32(3):230–236.
7. Elgendi M, Eskofier B, Dokos S, Abbott D. Revisiting QRS detection methodologies for portable, wearable, battery-operated, and wireless ECG systems. PLoS One. 2014;9(1):e84018.
8. Hannun AY, Rajpurkar P, Haghpanahi M, et al. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nat Med. 2019;25(1):65–69.
9. Ratasuk R, Vejlgaard B, Mangalvedhe N, Ghosh A. NB-IoT system for M2M communication. Proc IEEE Wireless Commun Netw Conf. 2016:1–5.
10. Lauridsen M, Krigslund R, Rohr M, Madueno G. An empirical NB-IoT power consumption model for battery lifetime estimation. Proc IEEE Veh Technol Conf. 2018:1–5.
11. Andres-Maldonado P, Ameigeiras P, Prados-Garzon J, Navarro-Ortiz J, Lopez-Soler JM. Narrowband IoT data transmission procedures for massive machine-type communications. IEEE Netw. 2017;31(6):8–15.
12. Martinez B, Monton M, Vilajosana I, Prades JD. The power of models: modeling power consumption for IoT devices. IEEE Sens J. 2015;15(10):5777–5789.
13. Shi W, Cao J, Zhang Q, Li Y, Xu L. Edge computing: vision and challenges. IEEE Internet Things J. 2016;3(5):637–646.
14. Zanella A, Bui N, Castellani A, Vangelista L, Zorzi M. Internet of Things for smart cities. IEEE Internet Things J. 2014;1(1):22–32.
15. Raza U, Kulkarni P, Sooriyabandara M. Low power wide area networks: an overview. IEEE Commun Surv Tutor. 2017;19(2):855–873.
16. Chen H, Li X, Zhao F. A lossy compression algorithm for wireless ECG monitoring based on compressed sensing. Sensors. 2015;15(3):5814–5835.
17. Sodhro AH, Sangaiah AK, Sodhro GH, Lohano S, Pirbhulal S. An energy-efficient algorithm for wearable electrocardiogram signal processing in ubiquitous healthcare applications. Sensors. 2018;18(3):923.